

Method and System for Optimizing Snow Flake Queries

BACKGROUND OF THE INVENTION

5

Technical Field

This invention relates to relational database systems. More specifically, the invention relates to optimizing snow flake queries.

Description Of The Prior Art

10

Relational database systems store large amounts of data, including business data that can be analyzed to support business decisions. Data records in a relational database management system in a computing system are maintained in tables, which are a collection of rows all having the same columns. Each column maintains information on a particular type of data for the data records which comprise the rows. Keys can be defined on a column or set of columns if the column or set of columns can uniquely identify each row of the table. In particular, one primary key can be defined on each table, which represents the primary identifier for each row of the table. A foreign key can be defined on one table which refers to the primary key of another table. A table can have multiple foreign keys defined, each of which refers to the primary key of a different table. Two tables can be joined together via the primary key - foreign key relationship, wherein this join will bring information from these two tables together.

25

Organizations are known to archive data in a data warehouse, which is a collection of data designed to support management decision making. Data warehouses contain a wide variety of data that present a coherent picture of business conditions at a single point in time. One data warehouse design implementation is known as a "star schema", also known as multidimensional modeling. The basic premise of star schema is that information is classified into two groups, facts and dimensions. A fact table

comprises the main database records concerning the organization's key transactions, such as sales data, purchase data, investment returns, etc. Dimensions are tables that maintain attributes about the data in the fact table. Each dimension table has a primary key column that corresponds to a foreign key column in the fact table.

5

The star schema provides a view of the database on dimension attributes that are useful for analysis purposes. This allows users to query on attributes in the dimension tables to locate records in the fact table. A query would qualify rows in the dimension tables that satisfy certain attributes or join conditions. The qualifying rows of the
10 dimension tables have primary keys that correspond to foreign keys in the fact table. A join operation is then specified to qualify rows of the fact table. Typically, the primary key columns of the dimension tables are compared against the corresponding foreign key columns in the fact table to produce the results.

15 Fig. 1 is a prior art diagram (10) of a star schema. The schema is "star like" with a central fact table (12), and peripheral dimension tables (22), (32), and (42). The central fact table (12) includes sales data, wherein each record includes information on the amount sold in the AMOUNT column (20), the time of the sale in the TIME_ID column (14), the product sold in the PROD_ID column (16), and the geographic region of the
20 sale in the NATION_ID column (18). The dimension tables (22), (32), and (42) provide attributes on the TIME_ID column (14), PROD_ID column (16), and NATION_ID column (18) in the fact table. The primary key columns of each of the dimension tables (22), (32), and (42) are the TIME_ID column (24), PROD_ID column (34), and NATION_ID column (44), respectively. The columns (14), (16), and (18) in the fact
25 table (12) are foreign keys that correspond to primary keys (24), (34), and (44) of the dimension tables (22), (32), and (42) that provide attributes on the data in the fact table (12). Accordingly, a join operation of the fact table (12) and the nation dimension table

(42) with the NATION_ID as the join key will bring necessary information from the nation dimension table (42) together with information from the fact table (12).

A "snow flake schema" is an extension of the star schema where there can be multiple levels of dimension tables. Fig. 2 is a prior art diagram (60) of a snow flake schema. As in the star schema of Fig. 1, there is a central fact table (12), and peripheral dimension tables (22), (32), and (42). In addition, there is a second level of dimension tables (62) and (72) which provide attributes on the MANU_ID column (36) of dimension table (32) and the REGION_ID column (46) of dimension table (42). A level for a table in a snow flake schema is defined as the number of join operations it needs to reach the fact table. As shown in Fig. 2, the fact table is at level 0, dimension tables (22), (32), and (42) are at level 1, and dimension tables (62) and (72) are at level 2. The primary key columns of the second level dimension tables (62), (72) are the MANU_ID column (64) and REGION_ID column (74), respectively. Column (36) in table (32) and column (46) in table (42) are foreign keys that correspond to primary keys (64) and (74) of dimensional tables (62) and (72), respectively. Accordingly, a join operation of the second level dimension tables (62) and (72), with the first level dimension tables (32) and (42), followed by a join operation with the fact table (12) will bring information from the two levels of dimension tables together with information from the fact table (12).

20

The fact table in a relational database is the largest table in the structure. As such, processing the fact table is resource intensive when compared to processing other tables. A query in a star schema or a snow flake schema usually has predicate filters on dimension tables and not on the fact table. A common technique for processing a query is to "push down" filter selectivity from dimension tables to the fact table. In a snow flake schema, filter selectivity from higher level dimension tables can be pushed down to the fact table via intermediate push down to lower level dimension tables. Predicate filters on a dimension table only select a subset of rows from the dimension table, and the

technique of push down enables processing of only rows from the fact table corresponding to the qualifying rows from the dimension table. However, one must balance the cost of conducting a push down and the benefit of this push down to determine whether a push down from this dimension table to the fact table will be
5 beneficial. The cost benefit analysis of conducting a push down depends on the selectivity of the predicate filter of the dimension table. A small selectivity indicates that a small number of rows will qualify based on this filter. Accordingly, the smaller the selectivity the dimension table has, the more beneficial it is to push down this filter selectivity to the fact table.

10

Much effort has been expended in developing optimization techniques for conducting queries in a snow flake schema. Fig. 3 is a prior art flow chart (100) illustrating a known method for a query in a snow flake schema. The initial step in the process is to input a query into the database (102). Following the input of the query into
15 the database, a test is conducted to determine if the query is a snow flake query (104). If the query is not a snow flake query, an alternative optimization technique will be conducted (106). However, if the query is a snow flake query, a level is assigned to each table in the query (108). The initial level for the query optimization process is set to one less than the maximum number of levels of tables in the query (110). Thereafter, for
20 each table at the current level the following steps are followed (112): the current table is assigned as the parent table (114), and all child dimension tables are ordered for the current parent table based on selectivity (116). Following the ordering of the child dimension tables, for each child dimension table of the current parent table the following steps are followed (118): a test is conducted to determine if the child dimension table
25 should be pushed down based on selectivity (120), a positive response to the test at step (120) will result in the marking of the child dimension table for push down (122), followed by a subsequent test to determine if the pushed down child dimension table is the last child dimension table of the current parent table (124). A negative response to

the test at step (124) will result in proceeding to the next child dimension table for the current parent table (126) followed by a return to step (120). However, a positive response to the test at step (124) as well as a negative response to the test at step (120) will exit the loop initiated at step (118) and will result in a subsequent test to determine if the parent table acted upon at step (114) is the last table at the specified level in the query (128). A negative response to the test at step (128) will proceed to the next table at the current level (130) followed by a return to step (114). However, a positive response to the test at step (128) will result in proceeding to the next lower level in the query structure (132). Thereafter, a test is conducted to determine if the next lower level in the query structure is less than zero (134), which is an indication that all tables in the query have been exhausted. A negative response to the test at step (134) will return to step (112), whereas a positive response to the test at step (134) will complete the query optimization process (136). Accordingly, the above outlined process for a query optimization in a snow flake schema initiates the query optimization at the peripheral nodes of the query structure and proceeds toward the central node, also known as the fact table.

As shown in Fig. 3, the prior art method for query optimization in a snow flake schema conducts the node selection process for a query at the peripheral nodes. However, this method is not efficient in the selection process in that it requires selection of each child dimension table on an individual basis. Accordingly, what is desirable is a time and resource efficient system and method for optimizing a query in a snow flake schema.

25

SUMMARY OF THE INVENTION

This invention comprises a method and system for optimizing a query in a relational database having a snow flake schema.

5

In one aspect of the invention, a method is provided for optimizing a snow flake query. A logical node is created, wherein the logical node is comprised of a child dimension table and all dimension tables rooted at the child dimension table. Thereafter, it is determined if the logical node can be committed for push down to a fact table.

10

In another aspect of the invention, a system is provided to optimize a snow flake query. A fact table and at least two child dimension tables are provided, as well as a logical node comprised of a child dimension table and all dimension tables rooted at the child dimension table. An optimization module is provided to pledge the logical node for push down to the fact table.

15

In yet another aspect of the invention, an article is provided with a computer-readable signal-bearing medium. Means in the medium are provided for storing data in a relational database having a fact table and at least two child dimension tables. Means in the medium are also provided for creating a logical node comprised of the child dimension table and all dimension tables rooted at the child dimension table, and means in the medium are provided for determining commitment of the logical node for push down to the fact table.

20

25

Other features and advantages of this invention will become apparent from the following detailed description of the presently preferred embodiment of the invention, taken in conjunction with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram of a prior art arrangement of database tables in a star schema.

FIG. 2 is a block diagram of a prior art arrangement of database tables in a snow
5 flake schema.

FIG. 3 is a flow chart illustrating a prior art query optimization process in a snow
flake schema.

FIG. 4 is a flow chart illustrating a query optimization process in a snow flake
schema according to the preferred embodiment of this invention, and is suggested for
10 printing on the first page of the issued patent.

FIG. 5 is a block diagram of an arrangement of database tables in a snow flake
schema with two logical nodes.

DESCRIPTION OF THE PREFERRED EMBODIMENT

Overview

In a star schema within a relational database, a fact table and multiple dimension
tables are provided with primary key columns and associated foreign key columns to
20 enable communication between the tables. A snow flake schema, as shown in Fig. 2, is a
known relational database structure of a plurality of dimension table levels. A logical
node is created to encompass an identified child dimension table, and all further
dimension tables that may be rooted at the identified child dimension table. The creation
of the logical node mitigates search space traversal as commitment of a lower level child
25 dimension table is inclusive of all dimension tables rooted at the committed lower level
child dimension table.

Technical Details

Fig. 4 is a flow chart (200) illustrating the process of optimizing a query in a relational database having a snow flake schema of a fact table and multiple levels of child dimension tables as shown in Fig. 2. The initial step in the process is to input a query into the database (202). Following the input of the query into the database, a test is conducted to determine if the query is a snow flake query (204). If the query is not a snow flake query, an alternative optimization technique will be conducted (206). However, if the query is a snow flake query, a level is assigned to each table in the query (208). Following assignment of levels to each table at step (208), a cumulative selectivity is calculated for all dimension tables at each level for all levels of the query (210). The cumulative selectivity calculation effectively creates a logical node comprised of a level 1 child dimension table, i.e. tables (22), (32) or (42) of Fig. 2, and all dimension tables rooted at this level 1 child dimension table, i.e. tables (62) or (72) of Fig. 2. The creation of these logical nodes effectively reduces a snow flake schema to a star schema, allowing algorithms for a star schema to be used at the logical node level. Accordingly, the first part of the optimization process includes calculating a cumulative selectivity for all dimension tables, leading to creation of logical nodes and subsequent reduction of a snow flake schema to a star schema.

Following the calculation at step (210) the level of the query optimization process is set to zero (212). As such, the query processing starts at the fact table. For each table in the current level (214) the parent table is assigned as the current table (216). Following the assignment at step (216), a test is conducted to determine if the parent table is marked as committed for push down to the fact table (218). A positive response to the test at step (218) will follow in a selection of a subsequent table at the same level (220), and skip the remainder of the processing steps for this parent table. However, a negative response to the test at step (218) will result in an ordering of all the child dimension tables for the current parent table based on the cumulative selectivity (222)

obtained at step (210). In an example where the current table is the same as the fact table, the response to the test at step (218) would always be negative. Alternatively, the test in step (218) could be skipped when the current table is the same as the fact table. Following the ordering of all child dimension tables for the current parent table at step

5 (222), each child dimension table must proceed through the following three steps (224): first, a test is conducted to determine whether the logical node rooted at this child dimension table should be pushed down to the fact table based upon the cumulative selectivity calculation (226), second, a positive response to the test at step (226) will result in all dimension tables in this logical node being marked as committed for push

10 down to the fact table (228), and third, a subsequent test is conducted to determine if the child dimension table in question is the last child dimension table of the current parent table (230). At step (226), the determination of push down includes consideration of previously committed logical nodes for push down to the fact table so that only an optimal quantity of logical nodes are committed for push down to the fact table. A

15 negative response to the test at step (230) will result in the query optimization process proceeding to the next child dimension table of the current parent table (232), followed by a return to step (226). However, a positive response to the test at step (230) as well as a negative response to the test at step (226) will complete the loop initiated at step (224). Thereafter, a subsequent test is conducted to determine if the current parent table is the

20 last table to be tested at the current assigned level of the query optimization process (234). A negative response to the test at step (234) will proceed to step (220) and the selection of the next table at the current assigned level of the query optimization process. However, a positive response to the test at step (234) will result in incrementing the assigned level in the query structure to continue the query optimization process at the

25 next level of dimension tables (236). Thereafter, a final test is conducted to determine if the incremented level from step (236) is greater than or equal to the maximum number of levels in the query structure (238). A negative response to the test at step (238) will result in a return to step (214), and a positive response will result in completion of the

query optimization process (240). Accordingly, the optimization technique shown in Fig. 4 proceeds through each level of the query structure starting at the central fact table and proceeding through each subsequent level of child dimension tables.

5 As previously mentioned, the creation of logical nodes effectively reduces a snow flake schema to a star schema. Therefore, all decisions as to whether to push down to the fact table are made in the process of iterating through level 1 child dimension tables. This is accomplished with the current level of the query optimization set to 0 in the optimization technique shown in Fig. 4, where the fact table is the parent table being
10 considered in the first iteration. Subsequent processing where the current level of the query optimization process is greater than zero is only necessary if partial push down is desired. A partial push down happens when filter selectivity from a higher level dimension table is pushed down to a lower level dimension table, but not to the fact table. Using the schema depicted in Fig. 2, an example of partial push down is when the filter
15 selectivity from table (62) is pushed down to table (32), but there is no push down from table (32) to the fact table (12). Accordingly, when a partial push down is considered, all dimension tables that have been marked as committed for push down to the fact table can be passed over.

20 As illustrated in the optimization technique shown in Fig. 4, the decision of whether to push down to the fact table is made for each logical node. Fig. 5 is a block diagram (300) illustrating the fact table (12), a first level dimension table (22), and two logical nodes (310) and (330). Logical node (310) is comprised of the first level dimension table (32) and the second level dimension table (62), while logical node (330)
25 is comprised of the first level dimension table (42) and the second level dimension table (72). The process of marking table (32) for push down to the fact table (12) incorporates all child dimension tables rooted at table (32), and in this example it is only table (62) as there are no other tables rooted at either tables (32) or (62). However, if there was a third

level child dimensional table rooted at table (62), this table would be incorporated into logical node (310) and marked for push down to the fact table (12). Similarly, the process of marking table (42) for push down to the fact table (12) incorporates all child dimension tables rooted at table (42), and in this example it is only table (72) as there are
5 no other tables rooted at either tables (42) or (72). The process of creating a logical node incorporates all decision making at the logical node level, *i.e.* at the dimension level of the lowest level table in the logical node. Once a logical node is committed for push down to the fact table, no further decision is required for any table within the logical node in subsequent query processing steps. Accordingly, the query optimization process
10 incorporates multiple levels of tables into a logical node for push down to the fact table.

Advantages Over The Prior Art

The process of creating a logical node that incorporates tables rooted at the selected child node mitigates search space traversal by the optimization module. The
15 calculation of cumulative selectivity that includes all selectivities from dimension tables rooted at the selected child dimension table enables the creation of a logical node. Once a logical node is committed for push down to the fact table, the search optimization process ceases consideration of any sub-tables related to the senior child dimension table for further push down considerations. Accordingly, the creation of a logical node
20 mitigates the search space required for traversal by the optimization module.

Alternative Embodiments

It will be appreciated that, although specific embodiments of the invention have been described herein for purposes of illustration, various modifications may be made
25 without departing from the spirit and scope of the invention. In particular, the concept of push down in the preferred embodiment is meant to include any method or mechanism for reduction of processing of the fact table based upon reduction of processing of dimension tables. In addition, the use of selectivity and cumulative selectivity in the

preferred embodiment includes any metrics used in the decision for push down to the fact table, either individually or in a collective manner. The concept of creation of a logical node is therefore not restricted to the calculation of cumulative selectivity, rather it includes any combination of metrics from all dimension tables within the logical node.

- 5 Accordingly, the scope of protection of this invention is limited only by the following claims and their equivalents.